9500420

9500420

30096

Estimating Requirements for Aircraft Recoverable Spares and Depot Repair

Executive Summary

John B. Abell

20101028 163

RAND | PROJECT AIR FORCE

R-4215-AF

Estimating Requirements for Aircraft Recoverable Spares and Depot Repair

Executive Summary

John B. Abell

Prepared for the
United States Air Force

RAND

PREFACE

This report summarizes the principal findings and recommendations of a body of research carried out in the project entitled, "Improving the Logistics Requirements Estimation Process." It discusses some important characteristics of the current system for estimating requirements for aircraft recoverable spares and depot repair and suggests several initiatives to improve data and data processing, requirements estimation, and operating policies.

Several other reports that describe the body of research summarized in this report are listed here:

- John B. Abell et al., *Estimating Requirements for Aircraft Recoverable Spares and Depot Repair*, RAND, R-4210-AF, 1993.

- John L. Adams, John B. Abell, and Karen E. Isaacson, *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts*, RAND, R-4211-AF/OSD, 1993.

- John B. Abell and Frederick W. Finnegan, *Data and Data Processing Issues in the Estimation of Requirements for Aircraft Recoverable Spares and Depot Repair*, RAND, MR-264-AF (forthcoming).

- Donald P. Gaver, Karen E. Isaacson, and John B. Abell, *Estimating Aircraft Recoverable Spares Requirements with Cannibalization of Designated Items*, RAND, R-4213-AF, 1993.

- Karen E. Isaacson and Patricia M. Boren, *Dyna-METRIC Version 6: An Advanced Capability Assessment Model*, RAND, R-4214-AF, 1993.

The first of these reports describes the entire body of work in considerably greater detail than this report and includes an elementary exposition of the current system. The second describes improved methods for forecasting the demand for aircraft recoverable spares and specifying the variance of the probability distribution describing the number of assets of a given type in resupply. The third report discusses data and data-processing issues related to the estimation of aircraft recoverable spares and repair requirements. The fourth presents a computational algorithm for estimating requirements for aircraft recoverable spares based on the assumption that items can be designated as cannibalizable or not. The fifth describes Dyna-METRIC Version 6, the capability assessment model used to evaluate

the stockage postures that were anticipated to eventuate from purchases of particular mixes of recoverable spares.

ACKNOWLEDGMENTS

CONTENTS

FIGURES

TABLES

1. BACKGROUND AND MOTIVATION

The research described in this report focuses on the difficult problem of estimating requirements for aircraft recoverable spares and depot-level component repair. These resources represent a major annual expenditure for the Air Force. In the mid-1980s, they absorbed roughly $5 billion annually, about $3 billion for spares investments and $2 billion for repair. The austere budgetary environment the Air Force currently faces reinforces the need to achieve adequate levels of aircraft availability at least cost. The research described in this report suggests initiatives that will enable the Air Force to reduce its investments in aircraft recoverable spares while maintaining roughly its traditional levels of aircraft availability. In short, these initiatives promise to enhance the robustness of stockage postures in the face of uncertain demands by producing a "smarter" mix of spares; they include fundamental changes in policy as well as improvements in requirements estimation techniques.

This work involved extensive analysis and evaluation of the current spares and repair requirements estimation system. Its principal thrust was to understand better the implications of management adaptations for spares requirements. By management adaptations we mean such initiatives as cannibalization; lateral supply; withdrawals of assets from war readiness spares kits (WRSK); and expedited repair, processing, handling, and transportation. These and other management initiatives enhance the performance of the logistics system but are not now accounted for in the computation of spares requirements. As we will show, uncertainty in demand degrades system performance, but management adaptations more than overcome that degradation, at least in peacetime operations. We have attempted here to try to quantify both effects.

In this report, we discuss our evaluations of the key initiatives to improve the spares and repair requirements estimation process and offer a set of recommendations that we believe will yield substantial improvements in the cost-effectiveness of spares and repair investments. We also discuss here our evaluations of an alternative approach to estimating the mix of depot-level component repairs. The reader is presumed to have some familiarity with the current spares and repair requirements estimation system. A more detailed and elementary exposition can be found in a companion report that describes many of the key features of the current system (Abell et al., 1993).

2. UNCERTAINTY AND MANAGEMENT ADAPTATIONS

The requirements database[1] changes rather substantially from year to year. New items appear, some items disappear, and our estimates of item characteristics change. Some of those changes are induced by errors of various kinds or by data problems, but most are real, i.e., they reflect evolving and changing field experience. We call the sum total of all such changes in the database *churn*. Churn costs money. When changes occur more or less randomly, some will induce increases in our estimates of item pipelines and some will induce decreases in those estimates. The problem is that when pipeline estimates decrease, the excesses thus induced are worth less than the cost of the shortages induced by increasing pipeline estimates. Thus churn plays an important role in shaping spares requirements and system performance, and it tends to contribute to the perception of long supply.

Figure 2.1 reflects the system's estimate of an item's pipeline based on eight quarters of past observation. The mean of the eight quarters is used as an estimator of the future characteristics of the item, including its pipeline. Thus, *ceteris paribus,* we expect to see eighty-odd assets of this type in resupply sometime in the future.

Unfortunately, the future seldom eventuates as we forecast. The F-16 avionics line-replaceable unit (LRU) represented by these data was no exception. As time passed, the failure rate of the LRU substantially improved. Given spares procurement actions consistent with the estimate reflected in Figure 2.1 and the actual item pipelines that evolved over time (shown in Figure 2.2), the system would have purchased roughly 60 assets more than needed.

Clearly, the change in the item pipeline could have been in the other direction and we might have bought too few assets. In general, this is the heart of the problem. We need to estimate our requirements for spares about 13 quarters in advance, on the average, and we are vulnerable to changing item characteristics, changing force structure, changing flying hour programs, and all of the other vagaries of the world that will affect the actual system performance that any

---

[1]We mean the D041 database, the set of data files maintained in the Recoverable Consumption Items Requirements System (D041).
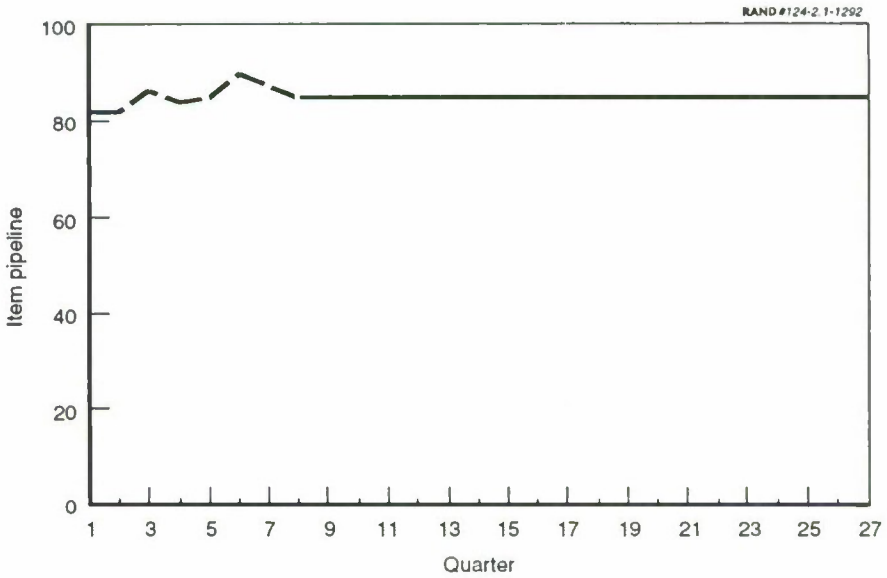
Figure 2.1—Estimate of an Item's Future Pipeline
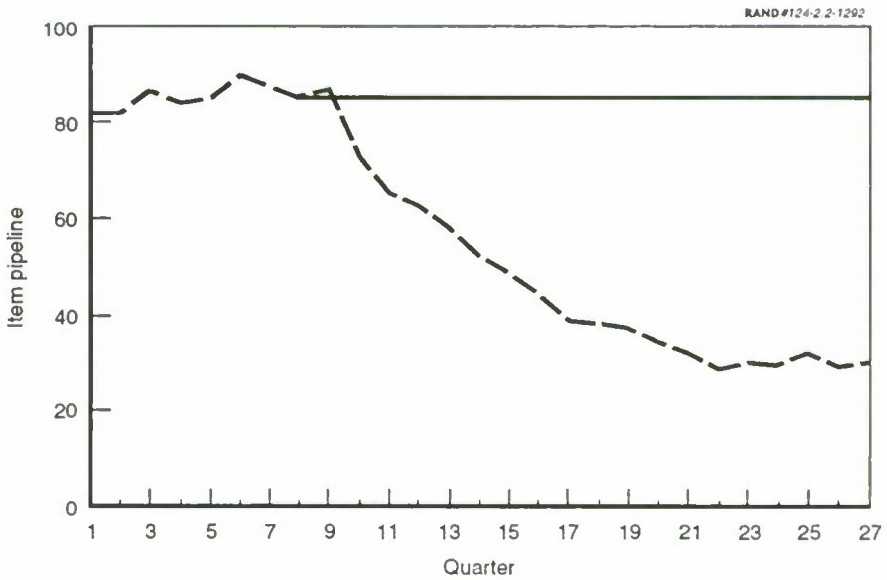
Figure 2.2—Evolution of an Item's Pipeline in Contrast
to the Estimate

specified mix of spares procurements will deliver. The estimation problem is made very difficult by such uncertainties.

Although churn and other kinds of uncertainties act to degrade system performance, management adaptations act in the opposite direction, more than overcoming the degradation resulting from uncertainty. We include in the category of management adaptations such actions as cannibalization (the consolidation of parts shortages in a minimal number of next-higher assemblies); prioritized processing, handling, and transportation of assets; mutual base support (lateral supply and repair); priority repair at both intermediate and depot level; and the withdrawal of serviceable assets from WRSK. The spares requirements system does not model most management adaptations, a characteristic that tends to induce overinvestment in spares. On the other hand it doesn't account for many of the uncertainties that pervade the requirements problem. On balance, the system tends to overinvest in aircraft recoverable spares, and it buys a less-effective mix of assets than it would if the initiatives discussed in this report were in place.

OUR APPROACH TO REQUIREMENTS EVALUATIONS

Figure 2.3 illustrates schematically our approach to the evaluation of the effects of churn and management adaptations on system performance given the current spares requirements system. We used the F-16 weapon system (all series) as a case study. The model of the requirements system includes the Aircraft Availability Model (AAM) used by the Air Force Materiel Command (AFMC) to compute requirements for safety stock (kindly provided to us by the Logistics Management Institute) and our replicas of the Central Secondary Item Stratification (CSIS) and Central Stock Leveling System (D028). This system of software enabled us to replicate the computations of the requirements system and allocate stock levels to bases and the depot as the Air Force actually does in practice.[2]

---

[2]An exception to the fidelity in our replication of the actual process is in our naive assumption that the requirement computed by the system is the actual mix of procurements. Item managers and others may intrude in this process in ways that change the mix of procurements recommended by the requirements system. Owing to lack of aircraft configuration data, we are forced to assume that items with application percentages less than 100 have the same application percentage at every location.
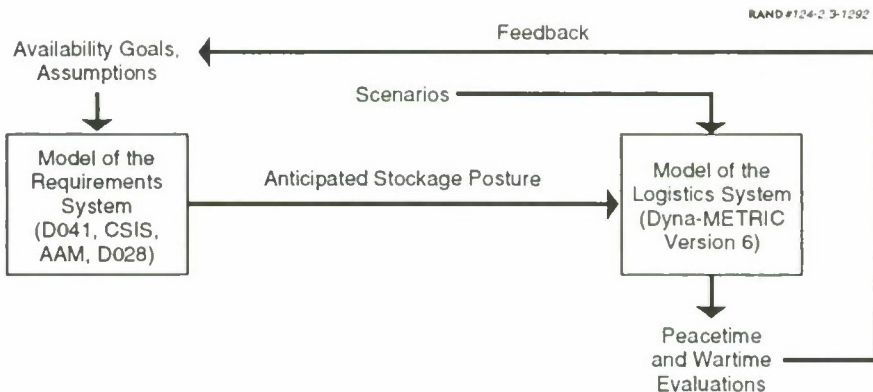
Figure 2.3—Our Approach to Requirements Evaluation Using the F-16 Weapon System as a Case Study

On the evaluation side, we primarily used the latest, most advanced hybrid analytic-simulation version of Dyna-METRIC, Version 6 (Isaacson and Boren, 1993), although we also used an earlier, purely analytic version, Version 4 (Isaacson et al., 1988). Version 6 enables us to evaluate the effects of management adaptations that we were not able to evaluate with earlier versions. It also explicitly models the indenture relationships among shop-replaceable units (SRUs), LRUs, and aircraft. Before discussing the results of the evaluations, it is important to understand the differences between the assumptions made in the requirements system's model of the world and those in Dyna-METRIC Version 6. Figure 2.4 helps point out the important differences.

Some of the differences between the requirements system and Dyna-METRIC derive from the definition of one key item characteristic: the number of users. The number of users of an item in the requirements system, the number used to determine the size of the pipeline for the average base, is the number of bases that had two or more demands for the item in the past 12 months. Thus its numerical value is a matter of chance. The number of bases in Dyna-METRIC is the number of bases at which the item will be exposed to demand, except for items that have an application percentage less than 100. In this case, Dyna-METRIC assumes that the same application percentage applies to every location. Since there is no convenient source of aircraft configuration data by tail number, this seems to be a reasonable assumption; however, it often overstates the number of bases at

which demands for items with application percentages less than 100 can occur.

The requirements system partitions item pipelines by dividing them by the (understated) number of users. This is referred to as the "average base assumption." Dyna-METRIC explicitly models the actual force beddown specified in its database. In this research we specified the force beddown by mission-design-series (MDS)-base, i.e., the number of F-16As, F-16Bs, Cs, and Ds at each location.

The requirements system assumes a steady-state system; Dyna-METRIC models dynamic scenarios, an especially useful feature for estimating wartime performance.

The requirements system assumes that parts shortages are randomly distributed among aircraft rather than being perfectly consolidated into the minimal number as does Dyna-METRIC Version 6.[3] Obviously, neither assumption is correct. Some parts are not readily cannibalized because of cost or risk. An assumption of perfect cannibalization is probably closer to the truth of actual practice than the assumption of random shortages; however, it leads Dyna-METRIC

RAND #124-2 4-1292

Requirements System	Dyna-METRIC Version 6
D041 number of users	Actual number of bases
Average base assumption	Explicit force beddown
Steady state	Dynamic scenarios
Random shortages	Cannibalization
No mutual base support	Lateral supply, repair
Past predicts future	Effects of churn
	Priority repair
	Other options

Figure 2.4—Dyna-METRIC Is Our Surrogate for the Real World

_____

[3]Dyna-METRIC Version 4, also used in many of our evaluations, provided us the option of full or no cannibalization. Since this research was completed, Dyna-METRIC Version 6 was modified to model full, partial (designated by stock number), and no cannibalization.

Version 6 to estimate system performance somewhat too optimistically.

There is no explicit assumption of lateral supply in the requirements system; however, it systematically understates the number of bases that may generate demands for an item. This understatement has an effect somewhat similar to an assumption of lateral supply, but the effect is not readily measurable. Dyna-METRIC explicitly assumes lateral supply as a user-specified option and enables the user to specify lateral repair, priority repair, and other options not of special interest in this discussion that make it an especially rich and powerful evaluative tool for logistics research and support of management decisionmaking.

The Quick Response Option

In the evaluations that follow we will refer to one management adaptation that we call the quick response option. We use the term to describe a more responsive logistics system in the sense that the system responds more quickly than suggested by the pipeline time segments reflected in the database. Although we model this responsiveness as applicable to all items, in actual practice one can achieve the same results by responding quickly with only a few of the most urgently needed items, those that tend to shape the performance of the system but which may not be known in advance. Table 2.1 reflects the assumption we made in modeling the quick response option. Note that actual flow days in the depot repair shop remain unchanged, but other segments of the total depot repair turnaround time are

Table 2.1

The Quick Response Option

	Pipeline Lengths	
	Actual	Quick
Base processing days	4[a]	3
Retrograde days	16	5 overseas, 2 CONUS
Supply to maintenance	3[a]	1
Shop flow days	40[a]	40
Service turn-in days	5[a]	2
Order and ship time	21	5 overseas, 2 CONUS
Total	89[a]	56 overseas, 50 CONUS

[a]Varies by item; average shown.

shortened by assumption. The total turnaround time is assumed to be reduced from an average of about 90 days to about 50 days. This shortening of the depot repair pipeline, such as might be achieved with an aggressive program of initiatives to enhance depot responsiveness, has powerful effects on system performance, as we will show.

THE EVALUATIONS

Figure 2.5 shows the result of our estimates of the peacetime performance of the current requirements system's anticipated stockage posture using the March 1986 D041 database. The evaluations were done using the same estimates of item characteristics (demand rates, not-repairable-this-station (NRTS) rates, repair times, etc.) reflected in the database used to compute the requirements. Thus no account was taken of the fact that when AFMC computes the requirement, it has no way of knowing how item characteristics will really eventuate three years or so in the future when the assets are received. Thus, these are "churn-free" estimates.
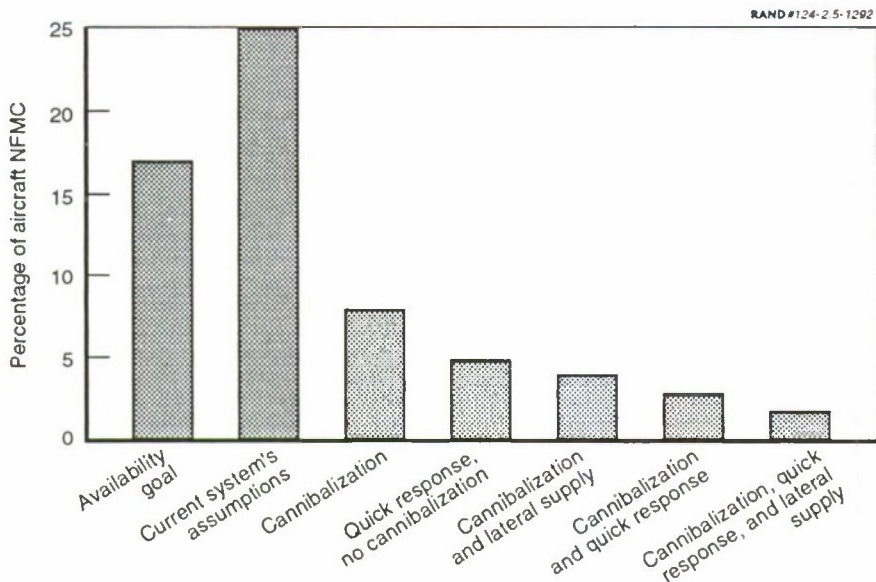

Figure 2.5—Peacetime Performance with WRSK Withdrawals

The leftmost bar in Figure 2.5 reflects the aircraft availability goal of 83 percent specified to the requirements computation, i.e., 17 percent of the force not fully mission capable owing to aircraft waiting for parts to be repaired or shipped to the location where needed. The bar labeled "Current system's assumptions" is Dyna-METRIC's estimate of the peacetime performance of the anticipated stockage posture using the assumptions of the requirements system. It is worse than the specified goal because of the assumption in Dyna-METRIC that items whose application percentages are less than 100 are used in equal proportion on the aircraft at every location, thus negating economies of scale in safety stock. As pointed out previously, there is no way to avoid this assumption given the data available. Dyna-METRIC Version 6 also assumes perfect consolidations of shortages, thus inducing an estimating error in the opposite direction. We do not claim that these errors are offsetting.

Note how dramatically management adaptations of the variety evaluated here improve system performance. Cannibalization or the quick response option alone achieves better than 90 percent availability. Lateral supply improves performance even more. The three adaptations together achieve 98 percent aircraft availability, far more than the specified goal. Recall, however, that these evaluations essentially assume perfect information in the sense that the same item characteristics are applied to the point in time in 1989 when this asset position is assumed to eventuate as were assumed in the requirements computation itself.

In Figure 2.6, the dark gray bars are the same as the bars in Figure 2.5: they represent the estimated performance of the anticipated asset position using the same item characteristics in the evaluations that were used in the requirements computation. The light gray bars in Figure 2.6 portray the estimated system performance using item characteristics that were observed at the time that the asset position would have evolved in the system, roughly three years after the spares requirements computation. Performance is significantly degraded owing only to our inability to forecast the future. This is what churn in the database really means: We will not achieve performance as good as we anticipate because we fail to account for the state-of-the-world and statistical uncertainties to which the system's performance is vulnerable.

If we made the same naive assumptions about the operation of the logistics system as are made in the current requirements computational system, we could expect to see the performance portrayed by the bar marked "Current system's assumptions." On the other hand, the
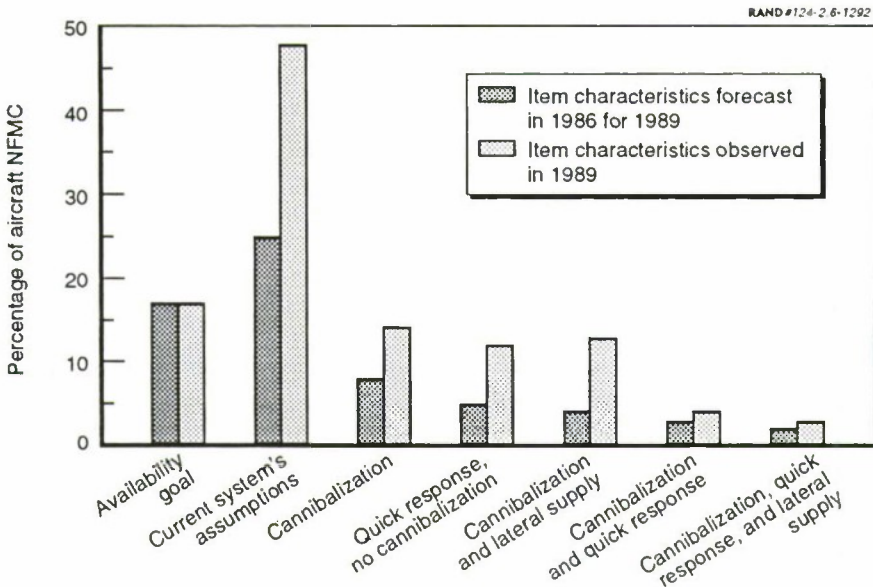
Figure 2.6—Effects of Churn on System Performance with
WRSK Withdrawals

logistics system doesn't really perform the way the requirements system assumes that it does since, in reality, management adaptations are in place and to some extent are routinely practiced even in peacetime. With such management adaptations systematically applied, the performance of the logistics system, even in the face of realistic levels of churn in the database, is dramatically improved, almost to the levels achieved assuming no churn at all.

Recall that Dyna-METRIC Version 6 assumes perfect consolidation of shortages of all items at a base, therefore overstating performance when cannibalization is involved. Actual performance is probably close to the case labeled "Cannibalization and lateral supply." F-16 not mission capable supply (NMCS) rates at the time of these observations were around 9 percent. In the face of such performance, availability goals of around 83 percent were persistently specified to the requirements computation. It is simply a numeric artifact of the aircraft availability model that such goal specification can be expected to produce much higher levels of availability in real life because management adaptations do have such powerful effects on system performance, especially in overcoming uncertainty.

In Figure 2.7, we add another set of bars, the white bars, to examine the effects of an alternative WRSK withdrawal policy, i.e., the WRSK is held inviolate during peacetime; no WRSK withdrawals are allowed. Note that this policy has little effect on system performance. Examination of the distribution of MICAP[4] termination codes suggests that WRSK withdrawals are frequently used during peacetime to satisfy MICAP requirements for serviceable spares. A WRSK withdrawal, even with the approvals required, is fairly quick and convenient. It obviates the need for the extra work associated with cannibalization and its attendant risk of damage and is quicker than a lateral supply action. If WRSK withdrawals are prohibited, however, lateral supply and cannibalizations would then be brought into play more frequently, thus avoiding serious degradation in performance, although transportation and cannibalization costs would clearly rise.
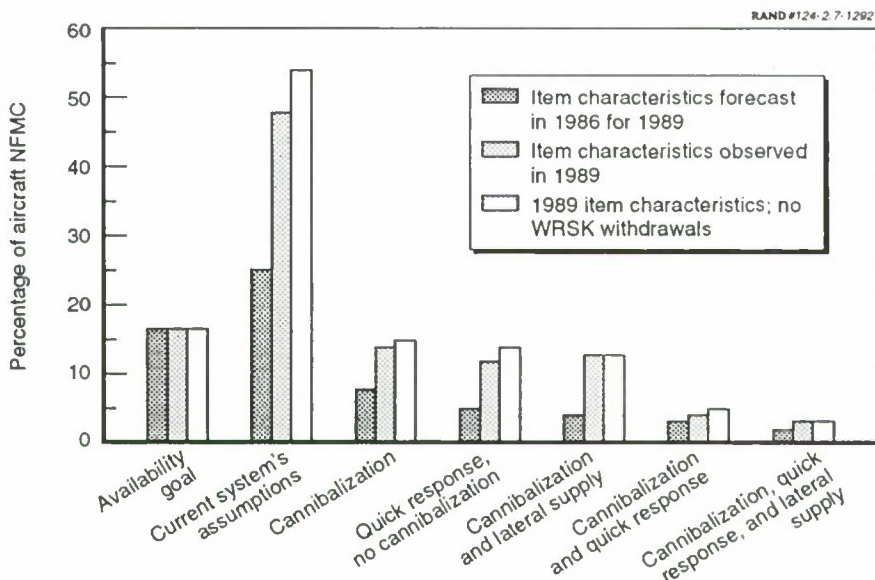

Figure 2.7—Effects of Churn on Peacetime Performance, No WRSK
Withdrawals Allowed

---

[4]Shortages judged to affect the mission capability of an aircraft.

Figure 2.8 reflects some assessments of wartime performance. The dark gray bars reflect the assumption of no churn; the light gray bars reflect the effects of churn. The measure of effectiveness is the proportion of aircraft not fully mission capable at the end of 30 days of war. The war is assumed to occur after 300 days of peacetime operation to ensure minimal effects of initial conditions. Full cannibalization is assumed in each of these cases. The current performance is roughly the same without churn as anticipated in the computation of the WRSK. However, accounting for churn shows that, without depot replenishment, we do not achieve the assumed performance.[5] With depot replenishment, quick response and lateral supply improve performance dramatically.

IMPLICATIONS OF THE EVALUATIONS FOR SPARES REQUIREMENTS

In summary, the evaluations show that management adaptations of the several varieties examined here more than overcome the degrading effects of the state-of-the-world and statistical uncertainties that make the spares requirements estimation problem so difficult. Many of the features of the current requirements system tend to induce a richness in spares investments. For example, the computational model, with traditional aircraft availability goals, buys safety stock to achieve levels of item availability that are virtually one.[6] Several other features of the current system result in the propensity for overinvestment. It seems clear that a leaner mix of stock levels in the system could easily be accommodated while maintaining desired levels of performance.

Conspicuous by its absence from the discussions in this section is the prioritization of repair at both base and depot level. Abell et al. (1992) showed that priority repair has very dramatic effects on aircraft availability when applied to depot-level component repair. Base-level priority repair also has important effects, and neither of those effects has been accounted for in these evaluations. Thus, in this sense at least, the evaluations are somewhat conservative, although the full cannibalization assumption of which we have included the effects is, as a practical matter, not achievable in the real

---

[5]The cases shown under "No depot" assume that units are cut off from depot resupply for the first 30 days of deployment. The "Depot" cases assume that depot resupply is not interrupted.

[6]Item availability is defined as the probability that an aircraft selected at random will *not* have a shortage of the item.
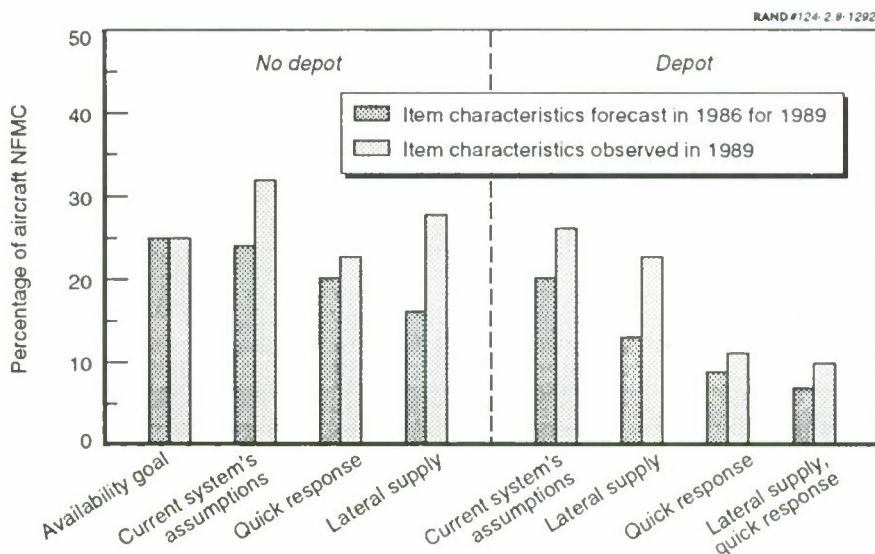
Figure 2.8—Estimated Wartime Performance,
WRSK + Primary Operating Stock

world. These two assumptions, full cannibalization and the omission of priority repair, act in opposite directions, but we do not have the data to quantify them.

The Air Force's goal should be to achieve satisfactory levels of aircraft availability at least cost, especially in the currently austere budgetary environment it faces. Thus the decision problem in spares require-ments estimation is to understand how to modify the requirements system so that substantial budgetary reductions can be achieved while maintaining acceptable levels of readiness and sustainability. Clearly, the requirements system should account for management adaptations; however, if it modeled *all* such adaptations, it would buy *too few* spares because it does not account for churn. The question is which adaptations should be accounted for, how, and to what extent, and what else should be done to improve the requirements process. In the remainder of this report we discuss some of the initiatives that we believe would contribute to the goal of achieving satisfactory levels of system performance at least cost and, to the extent possible, demonstrate their effectiveness.

3. KEY INITIATIVES TO IMPROVE THE PROCESS

The initiatives we discuss in this section fall into three categories: (a) data and data processing, (b) requirements estimation, and (c) policy initiatives. We discuss the initiatives in each category in turn. Unless otherwise noted, all examples are drawn from the March 1986 database.

DATA AND DATA-PROCESSING INITIATIVES

Random errors in the database have exactly the same kind of effect as churn does. They degrade system performance by inducing differences between estimated and actual factors and item characteristics. We encountered many data errors in the course of this work. We judge that two kinds of errors are especially important in their degrading effects, errors in item programs and errors in indenture relationships. We are also persuaded of the need for aircraft configuration and force beddown databases.

Errors in Item Programs

We were not able to identify the source of the errors in item programs. Item programs are reflected in the database. They are also easily computed using the application file, K004 (flying hour program) file, and straightforward arithmetic. Using the March 1989 database, we computed the item programs and compared them with the item programs entered in the database. We observed that 25 percent of them differed by 33 percent or more, an astonishing and clearly unacceptable level of error. We recommend that AFMC pursue the source of this error and correct the problem.

Errors in Indenture Relationships

The errors we observed in indenture relationships derive from the poor quality of the application data coupled with the practice of promoting items to the next-higher level of indenture when their next-higher assemblies cannot be found in the application file. This problem is described at length in a companion report (Abell and Finnegan, forthcoming). Consider an SRU with applications to an LRU (stock number) and aircraft. If the application file does not contain a record for the LRU, the SRU is promoted to LRU status in the construction

of the indenture file. The result of this promotion is that shortages of the SRU are assumed to hold an aircraft down rather than an LRU; therefore, the requirements system invests more heavily in the SRU than it otherwise would. There are large numbers of items in the category of SRUs and lower-indenture parts that are promoted to higher levels of indenture. In particular, engines and engine modules have not been in the database; therefore, engine components were promoted to LRU status and were overbought. A recent decision was made to stock-list engine modules and include them in the requirements database, but engines are still not planned to be included.

We suggest a procedure whereby all items that are promoted in the process of building the indenture file be identified to item managers or equipment specialists who should be required to correct the appropriate application data or interchangeability and substitutability data. We judge that this is an important issue, especially in estimating depot-level component repair requirements, although we do not have the data required to estimate the magnitude of its pecuniary effects. We also recommend the inclusion of engines as well as engine modules in the database.

The Need for Aircraft Configuration and Force Beddown Data

As we pointed out previously, the number of users in the current system is used to partition item pipelines. It is simply the number of bases with two or more demands in the past 12 months. It is used in lieu of data that would accurately identify an item's programmed exposure to demand at locations worldwide through the use of aircraft configuration, force beddown, and associated program data. The aircraft configuration database needed would contain a mapping of subgroup master stock numbers to aircraft tail numbers. The Air Force has purchased such configuration databases for its more recent weapon systems, but AFMC has not always provided for their updating. They could be routinely updated by data that reflect technical order compliances (TOCs) and implementations of engineering change proposals (ECPs) on specific tail numbers. All that would be needed is a mechanism for overlaying ECP and TOC data into the database. The data would need to include revisions to the mapping of subgroup master stock numbers to tail numbers implicit in the TOC or ECP. It is a straightforward data-processing problem.

A configuration database is needed in execution as well as in requirements estimation, for example in prioritizing the allocation of assets among locations worldwide to achieve specified aircraft availability goals. The number of users as currently defined understates

the true number and arguably also understates the effects of lateral supply. Use of aircraft configuration data, coupled with explicit accounting for lateral supply, could yield a more cost-effective mix of spares by taking explicit account of all locations at which items are exposed to organizational and intermediate maintenance (OIM) demand.

As we will discuss subsequently, such a database would also support consolidation of WRSK at a central location and yield substantial cost reductions in WRSK investment. It would also support the computation of base-specific application percentages for DRIVE[1] as well as the requirements system and would support the central allocation of stock levels to achieve specified aircraft availability goals without relying solely on demand-based criteria as D028 does.

INITIATIVES TO IMPROVE REQUIREMENTS ESTIMATION

The initiatives we suggest to improve the estimation of spares requirements include: (a) the assumption that designated items are cannibalizable while others are not because of cost or risk of damage, (b) the use of improved methods for forecasting demands and specifying the variance of the probability distribution of the number of items of each type in resupply (the pipeline distribution), and (c) correct specification of the number of users of an item and explicit representation of lateral supply.

Designated Cannibalization

AFMC recently moved to the use of a designated cannibalization policy for its computation of war readiness spares requirements. Designated cannibalization simply means that items are designated in the database as cannibalizable or not, or difficult or easy to cannibalize. The computation then treats them explicitly as cannibalizable or noncannibalizable. The result is a lower cost to achieve specified aircraft availability goals in contrast to the current system's assumption of randomly distributed shortages among aircraft. Traditionally, the Air Force has opposed, as a matter of policy, the assumption of cannibalization in its estimation of spares requirements. The current austerity of the budgetary environment, coupled with the ability to designate which items are cannibalizable and which are not, considerably weakens that position.

---

[1]A component repair prioritization and asset allocation mechanism, Distribution and Repair in Variable Environments.

We did not have data available that reflect which items are cannibal-izable and which are not. In an effort to understand the potential payoffs of a designated cannibalization policy, we arbitrarily desig-nated items by federal stock class as cannibalizable. In that process, we designated about 58 percent of the items in the database as canni-balizable. Although we do not claim any expertise in this process we have some reason to believe that the 58 percent number may not be too far from reality since about 70 percent of war readiness items were designated cannibalizable. WRSK items tend to be the higher-demand items and tend, too, to be more readily replaced, perhaps, than the typical item. The results of the exercise for the F-16 weapon system are reflected in Figure 3.1.

The payoffs are clear even without the precision we would have liked. In this case, computing spares requirements with an assumption of designated cannibalization reduces spares costs by about a quarter of a billion dollars for the whole inventory system (all aircraft) and, as suggested by the F-16 case, delivers roughly the same performance. The approach should be verified using maintenance technicians to designate the cannibalizable items. It seems like a very cost-effective policy. An approach to implementing it is discussed in Gaver et al. (1993).
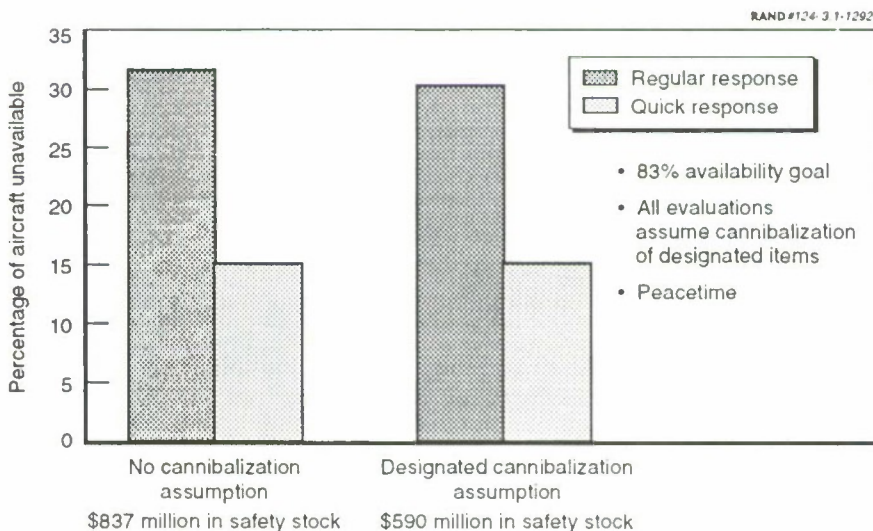

Figure 3.1—Assuming Designated Cannibalization Pays Off

Improved Demand Forecasting and Variance Specification

The current system uses an eight-quarter moving average to estimate item demand rates and NRTS rates. The variance of the probability distribution describing the number of items of each type in resupply is estimated using the observed variance-to-mean ratio (VTMR) of past demands. Demands are assumed to be strictly proportional to flying hours (for flying-hour-driven items), although we have never seen data that confirm this relationship. The assumption is made that demand processes are stationary, compound Poisson processes. They are, in general, not compound processes but, rather, nonstationary processes.[2] The current procedure of specifying the variance ignores forecasting uncertainty, accounting only for the stochastic variability in the demand process. It is approximately correct numerically for 13-quarter procurement lead times, but quite by chance.

The eight-quarter moving average is not sufficiently sensitive to nonstationarity in item demand processes. When an anomaly occurs, it takes two years before its effects vanish. Moreover, it gives as much weight to observations several quarters old as it does to very recent observations.

We explored a variety of alternative approaches to forecasting demands and specifying the variance of the pipeline distribution. The improved methods we chose, weighted regression forecasting and variance specification that explicitly estimates forecasting uncertainty as a function of the length of the planning horizon, reduce the demand forecasting error by roughly 40 to 50 percent on high-demand items (15 or more demands per quarter) over planning horizons of interest in the requirements problem. Table 3.1 shows the improvement in both root mean squared error (RMSE) and mean absolute deviation (MAD). The improved methods are described in detail in a companion report (Adams et al., 1993).

Table 3.1

Percentage Improvement in RMSE and MAD of Improved Techniques over Current System

Measure	10-Quarter Horizon	13-Quarter Horizon
RMSE	48	38
MAD	51	45

[2]These issues are discussed at greater length in Adams et al. (1993).

The use of the improved methods of forecasting and variance specification resulted in an estimated cost reduction of almost a quarter of a billion dollars ($239 million) in the requirements computation done with the March 1986 database while achieving a level of performance almost equal to that of the current system. Table 3.2 reflects the results of our evaluations of these methods with Dyna-METRIC Versions 4 and 6.[3] War readiness spares were not included in these evaluations.

It should be noted that forecasting errors degrade system performance in much the same way as data errors and churn do. They cause the system to anticipate conditions that do not eventuate. Such errors cost money, just as churn and data errors do, and they diminish the cost-effectiveness of our spares investments. Improvements in forecasting, like correction of data errors, help diminish the levels of uncertainty surrounding our resource allocation decisions.

Correct Specification of the Number of Users and Explicit Modeling of Lateral Supply

We are unable to assess the implications of moving to a system that incorporates the use of aircraft configuration and force beddown data simply because we have no way to estimate the effects of having correct base-specific application percentages for items whose true percentages are less than 100. In an attempt to understand this issue better, we made four evaluations with Dyna-METRIC Version 4, the results of which are shown in Figure 3.2. Each of the four evaluations assumed cannibalization but no lateral supply (since Version 4 cannot model lateral supply). Each used an aircraft availability goal of 83

Table 3.2

Performance of Improved Techniques Compared to Current System

Management Adaptations	Percentage of Aircraft Unavailable, Peacetime	
	Current System, $3709 Million	Improved Methods, $3470 Million
No cannibalization	74.9	76.3
Full cannibalization	33.0	33.1
Cannibalization, lateral supply	17.3	17.2
Cannibalization, quick, lateral supply	3.2	3.6

[3]Again, observed VTMRs of demands were used in these evaluations, capped at 5.0 as before.

percent. The variance-to-mean ratios in these evaluations were speci-
fied to be the same as those observed during the period April 1988
through March 1990, except that the observed values were capped at
5.0 to avoid being too vulnerable to data errors. The capping affected
about 5 percent of the items.

The leftmost bar in Figure 3.2 represents the current system's stock-
age posture that results from using the "number of users" data ele-
ment from the D041 database. The budget level was $3709 million,
as before. Using that budget level, we ran the requirements computa-
tion a second time using the correct number of users, i.e., the number
of bases at which items would be exposed to OIM demand in Dyna-
METRIC. Performance improved by 9 percent of the possessed air-
craft as shown in the second, lighter gray bar. Next we ran the re-
quirements system using the number of bases as the number of users.
Because increasing the number of users with the same item pipelines
gives up economies of scale in safety stock (since the safety stock
must be allocated to more locations), the budgetary requirement in-
creases. At this higher budget level, we then evaluated performance
using the number of users from D041 against the "correct" number of
users and, again, using the number of bases as the number of users
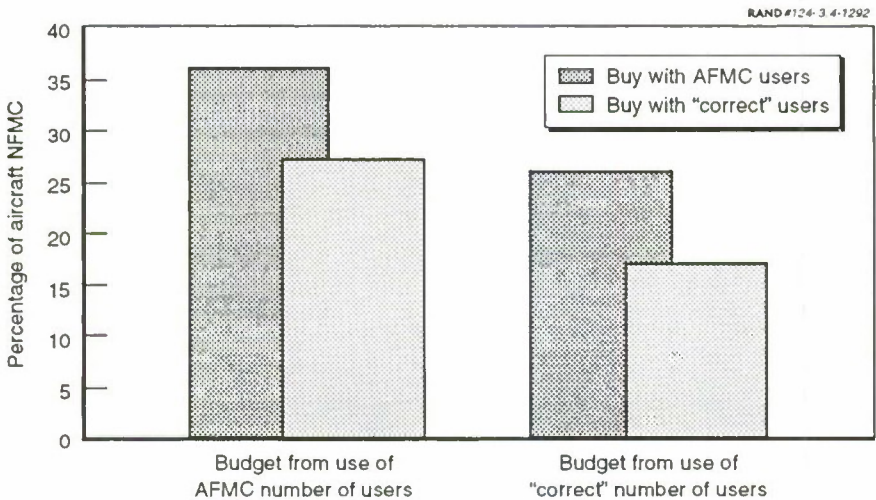improved performance by 9 percent of the possessed aircraft.

Figure 3.2—Effects of Consistency in the Number of Users

The results in Figure 3.2 are persuasive. For either budget level, getting the number of users "correct" in the sense that it is consistent between the requirements computation and the evaluation model (our surrogate for the real world) increases aircraft availability significantly.

The use of the actual number of bases at which an item is exposed to OIM demand (in lieu of the number of users in D041), the planned force beddown specified by aircraft tail number, and base-specific application percentages (in lieu of the average base assumption), coupled with an explicit model of lateral supply, would increase the cost-effectiveness of the anticipated stockage posture.

POLICY INITIATIVES

In the discussion that follows we explore two policy issues that have powerful implications for the cost-effectiveness of the spares requirements estimation process and for the logistics system at large. The first involves initiatives to improve the responsiveness of depot-level organic and contractual component repair. The second involves reevaluating the levels of war readiness spares required to meet the changing military threat and consolidating war readiness spares kits.

Improving the Responsiveness of Depot-Level Component Repair

As we have shown in the evaluations discussed in Section 2, quick response, i.e., reducing depot repair turnaround times, has powerful effects on logistics system performance and spares investments. Quick response is, of course, an abstraction for a more responsive logistics system. We have a definitive notion of responsiveness. We mean *relevant, timely, and robust*. By *relevant* we mean that the depot-level repair activity, whether contractual or organic, consistently repairs the components that are the most important to the achievement of aircraft availability goals. We have observed that the current system of negotiated quarterly repair goals and its associated performance measures and incentives does not produce a mix of assets especially relevant to the aircraft availability goals. *Timely* simply means that assets, especially the most relevant assets, stay in the depot repair pipeline the minimal time. *Robust* means that the repair facility is able to respond to urgent, unanticipated demands quickly so that its performance can be relied upon in the face of uncertain demands.

At the time of this writing, RAND is formulating research that is intended to explore and evaluate initiatives for enhancing depot-level

component repair responsiveness. We intend to explore issues related to: (a) congruence of goals between the depot-level repair activity and the combat force, (b) improved prioritization of retrograde, processing and handling, and transportation of serviceables, (c) improved material support for both contractual and organic repair, and (d) decision support systems for joint decisionmaking between operating commands and the air logistics centers in planning, programming, budgeting, and execution of depot-level repair, especially appropriate to the advent of stock funding of recoverables.

During the first two Coronet Deuce demonstrations, the air logistics centers involved were able to achieve significant reductions in component repair flow times through the depot repair cycle. This achievement was due in part to improved operating policies and procedures as well as motivation to make the two-levels concept work. We recommend that AFMC pursue such improvements aggressively for application to other workloads throughout the command and at contractors' facilities in a continuing program of initiatives to improve depot-level repair responsiveness.

Consolidating the Storage and Management of War Readiness Spares

An alternative policy for providing war readiness spares to deploying units is to maintain them centrally and develop the ability to determine and deploy the appropriate mix of spares with the deployment of any specified mix of aircraft. Current policy allocates WRSKs to every unit with a wartime deployment tasking. While this may have been consistent with a NATO scenario, it seems inconsistent with the changing global threat and the increased focus of attention on regional deployment contingencies.

The mix of spares in the WRSK is computed by WSMIS/REALM[4] and is adjusted during WRSK reviews conducted jointly by AFMC and the major commands to tailor the spares mix to the specific aircraft assigned to particular units. The need for such a process would be obviated by the aircraft configuration database described previously. The result of having a configuration database would be a reduced WRSK requirement. Rather than investing in a WRSK for each unit, economies of scale would make it possible to invest in enough war readiness spares to support contingency deployments, presumably

---

[4]WSMIS is AFMC's Weapon System Management Information System. REALM (Requirements Execution Availability Logistics Module) is the part of WSMIS that computes WRSK requirements.

substantially less than implied by a NATO scenario. Given a set of tail numbers specified for deployment, the computation of the appropriate spares mix could be done in a matter of minutes. The spares would then be pulled from storage and shipped to the deployed unit. The time required to assemble the required WRSK could be minimized by maintaining a generic kit by aircraft MDS and replacing only those items that did not apply to the particular tail numbers deployed. An exception listing could easily be produced by the computational system.

The F-16 again provides an interesting example. The current WRSK authorization for F-16 units worldwide, for all series, includes 50 individual WRSKs intended to support deployments of from 6 to 48 aircraft, plus 65 packages of war readiness spares to support the LANTIRN, ALQ-119, ALQ-131, ALQ-184, and QRC-80.01 systems. The total worth of these spares is about $604.1 million (based on the costs reflected in the current WRSK database maintained by Headquarters, AFMC). If one were to invest in enough spares to support a deployment twice the size of Operation Desert Shield/Storm (ODS) with 30 days worth of spares usage at 100 percent NRTS rates, the total requirement would be about $276.8 million, a reduction of $327.3 million, just for the F-16 weapon system.

With such a database, a system of consolidated WRSK storage, maintenance, computation, packaging, and deployment could become a reality and would yield significant personnel reductions at base level as well as substantial reductions in WRSK requirements. System performance without war readiness spares at bases would degrade somewhat, but not seriously, as we showed in Section 2 (Figure 2.7). It is not unreasonable to hypothesize that the correct specification of the number of users in the spares requirements computation could overcome this degradation but, of course, the correct number of users is unknown without the configuration database.

The Air Mobility Command has already moved to a centrally managed WRSK, apparently very successfully, but it enjoys an essentially homogeneous configuration among its airlift aircraft of a given MDS. In the case of the Tactical Air Forces, when the need for deployment of a specified number of aircraft arose, a WRSK could be assembled and deployed to the desired location. Unfortunately, if there is significant heterogeneity among individual aircraft of a given MDS, there is no way for the central manager to know what specific stock numbers to include in the WRSK without an aircraft configuration database.

Figure 3.3 reflects the investments associated with enough war readiness spares to support ODS and one-and-one-half and two times the ODS deployment for the F-16 aircraft. It also shows the investment required under the current policy of equipping every deployment-tasked squadron with its own WRSK as was done for the NATO scenario. A consolidated WRSK policy could, under current views about wartime contingency planning, deliver very attractive savings while not degrading *needed* combat capability. Moreover, its implementation could be coupled with a prudent WRSK withdrawal policy that would allow the withdrawal of some assets from the WRSK.
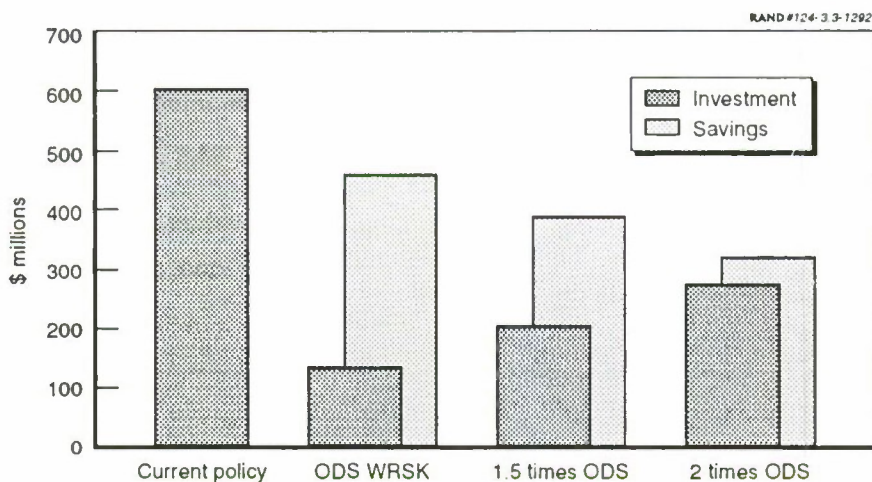

RAND #124-3.3-1292

Figure 3.3—Consolidating WRSK Yields Substantial Savings

4. ESTIMATING COMPONENT REPAIR REQUIREMENTS

In a prior RAND report, we described the prototype of a component repair prioritization and asset allocation mechanism called DRIVE. DRIVE was extended to apply to the problem of estimating quarterly component repair requirements. Its effectiveness in this application was described in Abell et al. (1992). A later evaluation of DRIVE in its role in estimating quarterly component repair requirements is included in a companion report to this one (Abell et al., 1993).

Both of these evaluations conclude unambiguously that, given accurate data, DRIVE is clearly superior to the current system in estimating quarterly repair requirements. We recommend AFMC continue its efforts to resolve the policy and implementation problems associated with commandwide implementation of the production version of DRIVE and use it in estimating quarterly repair requirements.

DRIVE's superiority to the current repair requirements estimation process lies in several features that differ markedly from the current system:

- DRIVE takes explicit account of a very recent "snapshot" of the worldwide asset position and provides means to display asset positions graphically in a very intuitively appealing and easily understandable manner.
- It is capable of explicitly identifying the catch-up requirement.
- It offers arguably better criteria for identifying critical items.
- It is dramatically more responsive to the current needs of the combat force than is the current system.

5. RECOMMENDATIONS AND SUMMARY

In this final section, we summarize our recommendations and provide a few, very brief closing remarks.

RECOMMENDATIONS

As before, our recommendations fall into three categories: (a) data and data processing, (b) requirements estimation, and (c) policy initiatives.

Data and Data-Processing Recommendations

Refer all items promoted during construction of the indenture file to item managers or equipment specialists so that their application data or interchangeability and substitutability data can be corrected.

Resolve the differences between D041 item programs and those computed from the K004 file and application file.

Develop an aircraft configuration database that maps subgroup master stock numbers into aircraft serial numbers and the means to maintain it systematically.

Recommendations to Improve Requirements Estimation

Implement a designated cannibalization assumption for POS requirements estimation.

Implement the improved demand forecasting and variance specification techniques described in Adams et al. (1993).

Include both engines and engine modules in the requirements database.

Evaluate the payoffs of defining the number of users correctly and explicitly modeling lateral supply.

Evaluate the use of base-specific data in requirements estimation, including application percentages, expected demands, and aircraft availability goals.

Estimate quarterly component repair requirements with DRIVE.

Policy Recommendations

Establish an aggressive and continuing program of initiatives to improve the responsiveness of depot-level component repair at both contractual and organic facilities.

Consolidate the storage and management of war readiness spares with the ability to compute, assemble, and deploy WRSKs in 24 hours with 6 hours' notice.

SUMMARY

In the requirements computations done for fiscal year 1987 (with the March 1986 database), these initiatives would have reduced spares procurement requirements by over $1 billion and delivered roughly the same aircraft availability. Their implementation would continue to yield savings each year, and they would help mitigate the perceived long supply problem.

BIBLIOGRAPHY

Abell, John B., et al., *DRIVE (Distribution and Repair in Variable Environments): Enhancing the Responsiveness of Depot Repair*, RAND, R-3888-AF, 1992.

Abell, John B., and Frederick W. Finnegan, *Data and Data Processing Issues in the Estimation of Requirements for Aircraft Recoverable Spares and Depot Repair*, RAND, MR-264-AF (forthcoming).

Abell, John B., et al., *Estimating Requirements for Aircraft Recoverable Spares and Depot Repair*, RAND, R-4210-AF, 1993.

Adams, John L., John B. Abell, and Karen E. Isaacson, *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts*, RAND, R-4211-AF/OSD, 1993.

Gaver, Donald P., Karen E. Isaacson, and John B. Abell, *Estimating Recoverable Spares Requirements with Cannibalization of Designated Items*, RAND, R-4213-AF, 1993.

Isaacson, Karen E., and Patricia Boren, *Dyna-METRIC Version 6: An Advanced Capability Assessment Model*, RAND, R-4214-AF, 1993.

Isaacson, Karen E., et al., *Dyna-METRIC Version 4: Modeling Worldwide Logistics Support of Aircraft Components*, RAND, R-3389-AF, May 1988.

O'Malley, T. J., *The Aircraft Availability Model: Conceptual Framework and Mathematics*, Logistics Management Institute, Task AF 201, Washington, D.C., June 1983.